

DECOMPOSING IMAGES INTO TRIANGLES BY DELAUNAY POINT PROCESSES

Dengfeng Chai*

Key Laboratory of Geoscience Big Data and Deep Resource of Zhejiang Province,
School of Earth Sciences, Zhejiang University, China - chaidf@zju.edu.cn

Commission II

KEY WORDS: Image partitioning, Delaunay point processes

ABSTRACT:

We propose a method for decomposing images into triangles. Contrary to superpixel methods, our output representation both preserves the geometric information disseminated in input images, and has an attractive storage capacity. Our method relies on the flexibility and efficiency of Delaunay point processes to address the problem. These stochastic models distribute points interacting between each other through Delaunay triangulations. The mechanism for distributing points combines several complementary ingredients including image discontinuity preservation, radiometric homogeneity inside atomic regions as well as priors on the shape of these regions. Said differently, sampled points and induced shapes work in tandem. The potential of our approach is shown through comparisons with existing oversegmentation methods and applications to vision problems.

1 INTRODUCTION

While traditional vision problems as image segmentation or stereo matching were mainly addressed at the pixel level in the last decades, an increasing number of works now proposes methods at the level of atomic regions. The preliminary decomposition of images into atomic regions is very successful as it allows for (i) an efficient way to preserve boundaries, and (ii) a better scalability than traditional pixel-based methods. The most common representation of atomic regions is superpixel (Ren and Malik, 2003), ie a set of connected pixels. If many superpixel methods have been proposed in the literature, few works have proposed alternative representations in which the geometric dimension of atomic regions is exploited. In this work, we propose a new representation for decomposing images into triangles. Despite their apparent simplicity, this geometric shape allow for highly flexible representation while insuring interesting geometric guarantees on the output decomposition.

The proposed representation is an alternative to superpixels that (i) preserves the geometric information disseminated in input images, and (ii) has a low storage capacity. More precisely, our objective is to design a geometric decomposition mechanism of images that conciliates both *efficiency* and *flexibility*.

To be efficient, the creation of keypoints and induced shapes must be performed simultaneously, and not sequentially. Sequential methods as (Tuylaars, 2010) can preserve boundaries correctly, but they have no control on region homogeneity. This observation motivates us to adopt a point sampling strategy in which the point distribution is guided by both boundary preservation and region homogeneity.

To be flexible, the geometric decomposition must be general enough to describe any type of scene. In particular, the point sampling mechanism must be able to generate configuration of points supporting free-form shapes. This requires (i) exploiting general shape priors, eg no restriction to line-segments (Duan and Lafarge, 2015), and (ii) having a point density high enough to capture details, contrary to (Ren et al., 2005) for instance. Triangles are a natural choice as unit surface elements for high flexibility.

* Corresponding author

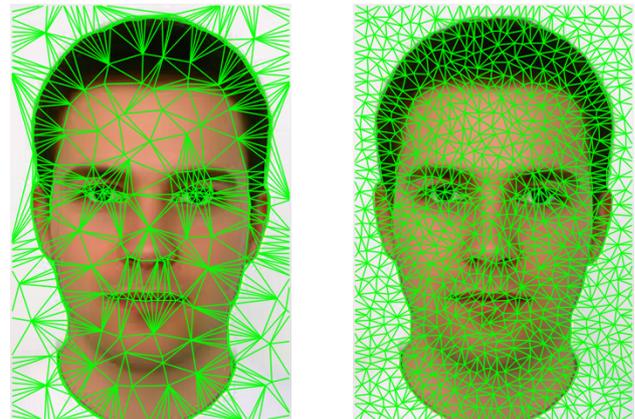


Figure 1. Image decomposition into triangles. Our Delaunay point process generates triangulations whose edges and facets capture the image contours and homogeneous areas respectively. The flexibility of the process allows for producing triangulations with low (left) or high (right) isotropy.

1.1 Contributions

Our algorithm takes an image as input and produces a decomposition into triangles of the image as output in which the image discontinuities are captured by edges while having homogeneous radiometry inside triangles as shown in Figure 1.

The key idea of our approach consists in guiding the point sampling by the induced geometric decomposition itself. For each sampled configuration of points, a Delaunay triangulation is used to decompose the image into triangles. Our algorithm brings three main contributions:

- **Delaunay point process.** We formulate our problem as a Delaunay point process (Bertin et al., 1999). This stochastic model, unexploited in Vision, is a spatial point process equipped of a Delaunay triangulation. It is particularly efficient as it allows points and induced shapes to work in tandem in a natural way.

- **Atomic decomposition into triangles.** We propose a geometric alternative to popular superpixels by decomposing images into triangles. In spite of its apparent simplicity, a triangulation offers flexibility as well as low storage capacity.
- **Shape control.** We propose a generic model in which the user can control the expected shape of regions. In particular, we develop shape priors for controlling the isotropy of triangulations as well as for capturing smooth, piecewise-smooth and linear image discontinuities.

After a brief introduction of related work in Section 2 and Delaunay point process in Section 3, our model formulation and the sampling procedure are detailed in Section 4 and Section 5. Experiments are presented in Section 6, followed by some applications to Vision problems in Section 7. Finally, we draw conclusions in Section 8.

2 RELATED WORK

Three main research directions are of interest in our review of previous work: image oversegmentation, extraction of keypoints, and shape reconstruction.

Image oversegmentation. The literature devoted to image partitioning into atomic regions is rich, in particular for superpixel generation. Existing methods, eg (Mori, 2005; Moore et al., 2008; Achanta et al., 2012; Van den Bergh et al., 2012; Liu et al., 2011; Veksler et al., 2010; Zhang et al., 2011) differ in terms of (i) methodology, eg region refinement or energy minimization on graph, (ii) output characteristics, eg boundary adherence, undersegmentation error, or region compactness, and (iii) performances, eg running times and memory consumption. They typically have complementary advantages, eg a high region compactness (Moore et al., 2008; Achanta et al., 2012) or a high boundary adherence (Liu et al., 2011; Van den Bergh et al., 2012). Partitioning images into geometric shapes offers several interesting advantages compared to superpixel decompositions, in particular a lower representation complexity and a better interpretation of the geometric information disseminated within the scene. Such methods are however more marginal in the literature. Delaunay triangulations or Voronoi diagrams (Aurenhammer et al., 2013) constitutes interesting geometric tools to partition images. In (Duan and Lafarge, 2015) for instance, images are partitioned into convex polygons. Although this method provides good results for man-made scenes composed of linear structures, it lacks of flexibility for more general scenes in which boundaries cannot be well approximated by line-segments.

Keypoint extraction. Partitioning into geometric regions usually relies on the placement of specific points that control the shape of regions. These points can either be detected through local descriptors (Mikolajczyk and Schmid, 2005) or sampled using stochastic processes (Descombes, 2011). The former option requires a sequential mechanism: keypoints are first detected, and then associated into shapes, eg 3D meshes (Labatut et al., 2007), and triangulation/quadrangulation (Tuytelaars, 2010). Outputs produced by these methods are strongly dependent of the keypoint detection. The later option combines simultaneously point positioning and shape coherence with respect to input images. Although this solution is more natural and mathematically elegant, extracted shapes are quite simple, mostly basic parametric shapes defined by a low number of parameters as circles or

rectangles, or low complexity planar graphs (Chai et al., 2013).

Shape reconstruction. Shape reconstruction methods often constitute a subsequent step to the atomic region partitioning problem. A natural way to partition the space with a geometric decomposition is to use Delaunay triangulation. This tool is exploited for instance for extracting object contours as polyline curves (de Goes et al., 2011) and for reconstructing object surfaces from MultiView Stereo (MVS) images as 2-manifold meshes (Labatut et al., 2007). Delaunay triangulations is also exploited for completing shapes using a Conditional Random Field formulation (Ren et al., 2005). Shape reconstruction can also be addressed by merging geometric entities as line-segments (Levinshtein et al., 2010; Sun et al., 2014).

3 DELAUNAY POINT PROCESSES

Delaunay point processes are mathematical models relying on both stochastic geometry and computational geometry. A brief introduction on these mathematical tools is given in the following.

Delaunay triangulation. Delaunay triangulation is a computational geometry tool for subdividing a planar domain into triangles (Lee and Schachter, 1980). Denoted by $DT(P)$, the Delaunay triangulation of a set of points P is defined such that no point in P is inside the circumcircle of any triangle of $DT(P)$. This space partitioning scheme has the interesting property to maximize the minimum angle in all the generated triangles. Performing atomic operations on such triangulations, eg removing or adding a point in P is also easy and computationally efficient.

Spatial point processes. These stochastic models describe an unordered set of points in a compact set $F \subset R^k$, where R^k is a k -dimensional space (Baddeley and Lieshout, 1993). For $n = 1, 2, \dots$, let Ω_n be the set of configurations $P = \{p_1, \dots, p_n\}$ that consists of n unordered points $p_i \in F$. A spatial point process on F is a mapping Ψ from a probability space to the set of configurations $\Omega = \bigcup_{n=1}^{\infty} \Omega_n$, such that, for all bounded Borel sets $S \subset F$, the number of points $N_{\Psi}(S)$ falling in S is a finite random variable. What makes spatial point processes attractive is the possibility of guiding the distribution of points by a probability density h , often expressed by a Gibbs energy $U \propto -\log h$. In this paper, we are interested in the point processes in dimension two, ie $k = 2$.

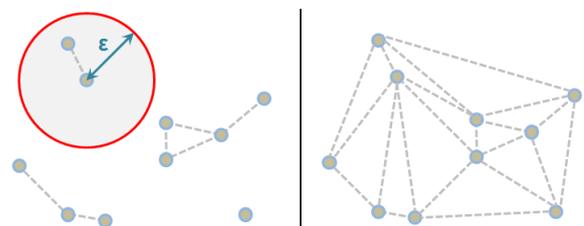


Figure 2. Spatial point processes. Spatial interactions between points (see dashed grey lines) in standard point processes are typically defined via a tolerance distance ϵ (left). Instead, Delaunay point processes (right) define interactions through Delaunay edges in a parameter-free way.

Delaunay point processes. Delaunay point processes are spatial point processes introduced by (Bertin et al., 1999) in which the

probability density h exploits Delaunay triangulations to define interactions existing between points. Contrary to standard point processes, Delaunay point processes guarantee that (i) each point interacts with at least two other points, and (ii) all points are connected within a unique interaction graph given by their Delaunay triangulation, also called Delaunay graph (see Figure 2). The Gibbs energy U that specified a Delaunay point process is defined as the sum of Delaunay-clique potentials:

$$U(P) = \sum_{\psi \in C(P)} g(\psi) \quad (1)$$

where $C(P) = C_1(P) \cup C_2(P) \cup C_3(P)$ is the set of cliques of cardinal one, two and three in the Delaunay graph. ψ is thus a point $p \in C_1(P)$, an edge $e \in C_2(P)$ or a triangular facet $f \in C_3(P)$ of the Delaunay triangulation, and $g(\psi)$ is its corresponding interaction potential.

4 MODEL FORMULATION

In this section, we specify the form of the energy U of the Delaunay point process we use to decompose an input image into triangles.

General form of the energy. Two types of information are exploited in the energy: the image consistency that measures the coherence between the triangle decomposition and the input image (ie a data term), and the shape prior that favors certain forms of triangle decompositions. Considering an energy that conforms to Equation 1, three types of geometric entities can be used, ie points, edges and triangular facets, to express image consistency and shape prior. Our energy is then formulated as

$$U(P) = \sum_{p \in C_1(P)} D_1(p) + \sum_{e \in C_2(P)} D_2(e) + \sum_{f \in C_3(P)} S(f) \quad (2)$$

where, D_1 and D_2 compose the image consistency term, and S corresponds to the shape prior.

Image consistency. To match triangles on the input image, we check that both points and edges are located on high image gradients. For efficiency reasons, we do not measure the radiometric homogeneity inside each triangle. In practice, such an additional term strongly increases running times for barely better results. The terms D_1 and D_2 are given by:

$$D_1(p) = -\log \|\nabla I(p)\| \quad (3)$$

$$D_2(e) = \frac{1}{|e|} \int_{p \in e} D_1(p) dp \quad (4)$$

where $\nabla I(p)$ is the image gradient at point p , and $|e|$ is the length of edge e . $D_2(e)$ equals to $D_1(p)$ averaged over all points p on the edge e . Note that more complex metrics for measuring image discontinuities, as (Arbellez et al., 2011), can also be used. Such a choice typically improves accuracy of results while increasing running times.

Shape prior. This term allows us to bring control on the shape of the triangles. Given the targeted application, one can expect different geometric characteristics on the output decomposition in terms of isotropy and regularity. As illustrated in Section 7, skinny triangles are more suitable to object polygonalization for instance as this application requires to distribute points densely

along the object contours and sparsely elsewhere. On the contrary, isotropic triangulations give better results for image abstraction for which atomic regions must rather be homogeneously distributed in the image.

As the shape of a triangle can be specified by the length of its edges, we propose a general shape prior of the form:

$$S(f) = \beta \sum_{i=0}^2 \min(\epsilon^2, (|e_i| - \mu_i)^2 / \sigma_i^2) \quad (5)$$

where $|e_0| \leq |e_1| \leq |e_2|$ represent the length of the three edges of facet f , the pairs (μ_i, σ_i) are the expected length and length variances of edge e_i to be specified by the user, β weights the importance of the shape prior with respect to the image consistency term, and ϵ is a minimum distance insuring points to not overlap. In our experiments, ϵ is fixed to 3. Figure 3 shows the behavior of the shape prior in three different cases.

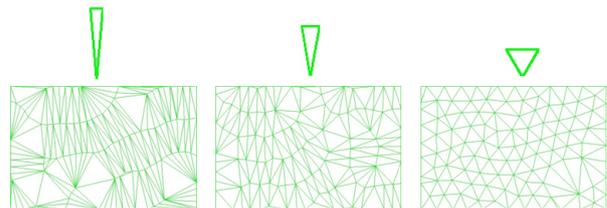


Figure 3. Shape prior. Three triangulations are simulated using different triangle shapes from low (left) to high (right) isotropy. Only the shape prior term S is considered in the energy U for this experiment.

5 SAMPLING

Monte Carlo sampling is traditionally used to explore the configuration space of spatial point processes and reach configurations close to the global minimum of U .

Algorithm 1 Sampling procedure

- 1- Distribute n points randomly on the image domain;
- 2- At iteration t under the current configuration $X_t = P$,

- Perturb a point $p \in P$ to a new configuration P' .
- Compute the Metropolis-Hastings ratio

$$R = \exp \frac{U(P) - U(P')}{T_t} \quad (6)$$

- Choose $X_{t+1} = P'$ with probability $\min(1, R)$, and $X_{t+1} = P$ otherwise
- Update the temperature by $T_{t+1} = C \cdot T_t$

Metropolis-Hastings. The Metropolis-Hastings algorithm (Green, 1995; Hastings, 1970) is an updating mechanism based on Monte Carlo. It consists in simulating a discrete Markov Chain $(X_t)_{t \in \mathbb{N}}$ on the configuration space Ω , converging towards an invariant measure specified by the Gibbs energy U . At each iteration, the current configuration of points P is perturbed according to proposition kernels into a new configuration P' . The perturbation is local, which means that P and P' are very close, and differ by only a few points. The configuration P' is accepted for the next iteration according to a probability depending on the energy

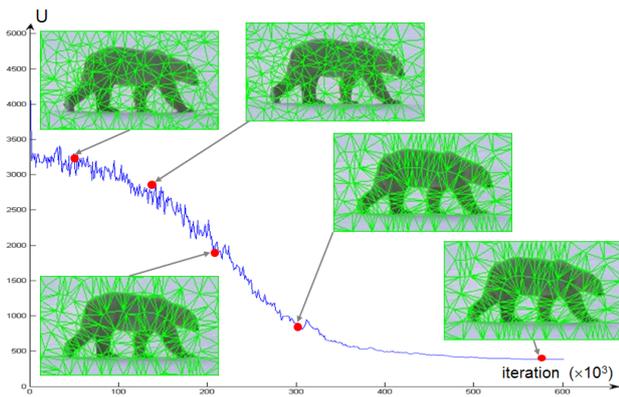


Figure 4. Energy evolution during the sampling procedure. Starting from a random initialization, the current configuration of points progressively evolves towards the optimal one with global minimum of the energy U .

variation between P and P' , the ratio of proposition kernels with respect to the reverse perturbation (in our case, this ratio is fixed to 1 as the proposition kernels are uniform distributions), and a relaxation parameter T that geometrically decreases at a rate C .

In practice, the sampling procedure is simplified by assuming the number of points is a constant value fixed by the user. This assumption is commonly exploited in oversegmentation methods for which the number of atomic regions is a model parameter. The sampling procedure is detailed in Algorithm 1.

Figure 4 shows the evolution of configurations and the energy decrease during the sampling procedure. In the first iterations, the current configuration of points is of bad quality, represented by a high energy. As the relaxation parameter decreases, the sampler progressively becomes selective and moves the current configuration towards the optimal solution.

Implementation. We implemented our algorithm in C++, using the Computational Geometry Algorithms Library¹ for geometric data structures and the associated operations. The manipulation of Delaunay triangulations is well adapted to the sampling procedure because the standard operations are (i) local similarly to the perturbations in the Metropolis-Hastings algorithm, and (ii) computationally efficient. In practice, when a point is affected by a perturbation, we distinguish two cases. If the point translation is short and does not modify the Delaunay graph, we only update the new coordinates of the considered point in the data structure. When the translation is large and affects the Delaunay graph, we first remove the considered point and then add the point with its new coordinates in the Delaunay triangulation.

6 EXPERIMENTS

We next describe a series of experiments to evaluate the flexibility and the efficiency of our algorithm, and to compare its performance to existing oversegmentation methods.

Model parameters. Our model formulation has several parameters. The most important parameters are the pairs (μ_i, σ_i) . These parameters allow us to specify the targeted shape of triangles, but also their size, and thus indirectly the expected number of triangles into the partition. We also use these values to estimate

¹ www.cgal.org

the number of points n to sample. Figure 5 illustrates the impact of the pairs (μ_i, σ_i) on output partitions. The weight β between the shape prior and the image consistency term is constant in our experiments and set to 1.

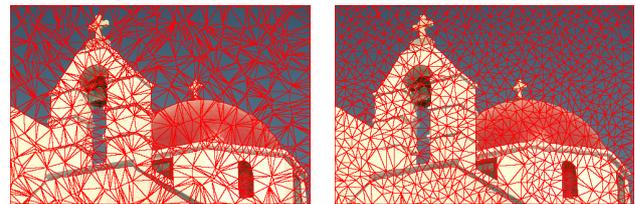


Figure 5. Impact of the shape prior. Selecting a long narrow triangle as targeted shape gives a low complexity and anisotropic triangulation (left). To the contrary, a short equilateral triangle as target generates a dense isotropic triangulation (right). The former is more suited to object contouring problems whereas the later is better for segmentation problems.

Flexibility. Algorithms that exploits geometric representations for images are usually poorly flexible and work only on specific types of scenes, eg man-made environments (Duan and Lafarge, 2015). As illustrated on Figure 7, the use of triangles allows us to capture free-form boundaries with good accuracy, in particular from natural object images, ie animals, faces or natural landscapes. The algorithm also performs well from man-made object images, eg the indoor scene, even if it does not preserve the linear structures as well as (Duan and Lafarge, 2015). Beyond the use of triangles, the possibility offered to the user to define a desired triangular shape allows him to tackle multiple applicative problems as developed in Section 7, as well as to insert complementary geometric priors as illustrated on Figure 6.



Figure 6. Point distribution simulated with an alternative geometric prior. Our shape prior can be completed by an additional term that favors smooth (left) and piecewise-smooth (right) angles between successive Delaunay edges. Points are then distributed regularly along curves. This complementary prior can be used for extracting free-form object contours.

Efficiency. Decompositions into triangles allow us to discover and preserve geometric information disseminated in images. In Figure 5 for instance, the contours of the church is captured by a subset of Delaunay edges. This provides valuable information related to the structure of the building as its shape, its orientation or its level of regularity. One of the main advantages of our representation is capacity storage. Contrary to superpixel methods that require to index each pixel of an image to define the regions, our representation can be saved in a very compact way as a triangulation graph. Another advantage is the memory consumption which is very low given the fact (i) the output representation is a low complexity graph, and (ii) the sampling procedure is memoryless. In terms of running time, our method requires a few seconds on images from the Berkeley database (Martin et

al., 2001), ie approximatively 150Kpixel images. This timing is reasonable for an energy minimization based approach. Note that our implementation is not optimized. In particular, running times could be strongly reduced by performing perturbations in parallel under GPU.

Comparisons with oversegmentation methods. Our algorithm has been evaluated using the standard quality measures for oversegmentation methods, ie boundary recall and undersegmentation error, and compare to two standard superpixel methods, ie SEEDS (Van den Bergh et al., 2012) and SLIC (Achanta et al., 2012), and a geometric method based on Voronoi partitioning (Duan and Lafarge, 2015). Figure 8 shows the comparison results from the Berkeley database (Martin et al., 2001). Superpixel methods (Achanta et al., 2012; Van den Bergh et al., 2012) perform better on both boundary recall and undersegmentation error. This result was expected as triangles can hardly compete with chains of pixels for capturing object boundaries. This is the price to pay by our method to have a simple and compact geometry representation. Our algorithm is more competitive against the Voronoi-based method (Duan and Lafarge, 2015), in particular on boundary recall for which Voronoi partitions are not flexible enough to capture accurately boundaries of natural objects. Note that region compactness is also usually considered as quality measure. Such a measure based on the regularity of pixel distributions with respect to the superpixel centroid is not adapted to triangular shapes and naturally scores low with our method.

Limitations. Our model formulation does not consider radiometric homogeneity inside each triangle. The quality of our results could be improved by introducing relevant metrics for which running times would not be significantly affected. Timings are reasonable for an energy minimization based approach, but remain not very competitive with respect to the fastest superpixel methods. Another limitation of our approach is its inability to combine decomposition and segmentation problems jointly.

7 APPLICATIONS

We present in the following a couple of applications showing the potential of our method for solving image and vision problems.

Polygonal contouring. Our algorithm can be used as a preliminary step to extract the contours of a specific class of objects. By assuming object contours correspond to cycles in the Delaunay graph, one can capture objects as polygons formed by Delaunay edges.

In Figure 9, building roofs are extracted from aerial images through polygons. In these examples, we first regroup connected triangles with similar radiometry by region growing. We then track edges on the border of the cluster of triangles to define the polygonal contours of roofs. While this procedure is very simplistic, most of roofs are correctly detected. In contrast, existing polygonal contouring methods, eg (Sun et al., 2014), rather exploit complex geometric primitive merging strategies.

Image abstraction. Our decomposition into triangles can also be used for instance in non-photorealistic rendering. As triangles capture homogeneous atomic regions of an image, one can simply color each triangle by the mean intensity of inside pixels to create an artistic representation of this image. As illustrated in Fig. 10, such vectorized representations are extremely compact and preserve nicely the details contained in original images.

8 CONCLUSION

We presented a method to decompose an image into triangles. Such a geometric partition contrasts with traditional superpixel representations whose storage capacity is often a limiting factor. The main originality of this work lies on the Delaunay point process we designed to partition images. Unexploited in Vision, we demonstrated that such a stochastic model is particularly interesting in terms of flexibility and efficiency, and for addressing some applications as polygonal contouring of objects or image abstraction.

In its current state, our algorithm cannot combine decomposition and segmentation jointly. In future works, we would like to investigate on the design of Delaunay point processes that can exceed this limitation. Excluding timing issues, one solution could be to create a two-layer point process with one extra layer devoted to the labeling of triangles. We also plan to extend our formulation in 3D to reconstruct surfaces from MVS images without being dependent of a preliminary keypoint detection step.

ACKNOWLEDGEMENTS

This work was supported by the National Natural Science Foundation of China (No.41571335). This work was supported by the Research Funding of School of Earth Sciences at Zhejiang University.

References

- Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P. and Susstrunk, S., 2012. Slic superpixels compared to state-of-the-art superpixel methods. *Ieee Transactions on Pattern Analysis And Machine Intelligence* 34(11), pp. 2274–2281.
- Arbelez, P., Maire, M., Fowlkes, C. and Malik, J., 2011. Contour detection and hierarchical image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33(5), pp. 898–916.
- Aurenhammer, F., Klein, R. and Lee, D.-T., 2013. *Voronoi Diagrams and Delaunay Triangulations*. World Scientific Publishing.
- Baddeley, A. and Lieshout, M. V., 1993. Stochastic geometry models in high-level vision. *Journal of Applied Statistics* 20(5-6), pp. 231–256.
- Bertin, E., Billiot, J. and Drouilhet, R., 1999. Spatial delaunay gibbs point processes. *Communications in Statistics. Stochastic Models* 15(2), pp. 181–199.
- Chai, D., Forstner, W. and Lafarge, F., 2013. Recovering line-networks in images by junction-point processes. In: *26th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2013, June 23, 2013 - June 28, 2013*, pp. 1894–1901.
- de Goes, F., Cohen-Steiner, D., Alliez, P. and Desbrun, M., 2011. An optimal transport approach to robust reconstruction and simplification of 2d shapes. *Computer Graphics Forum* 30(5), pp. 1593–1602.

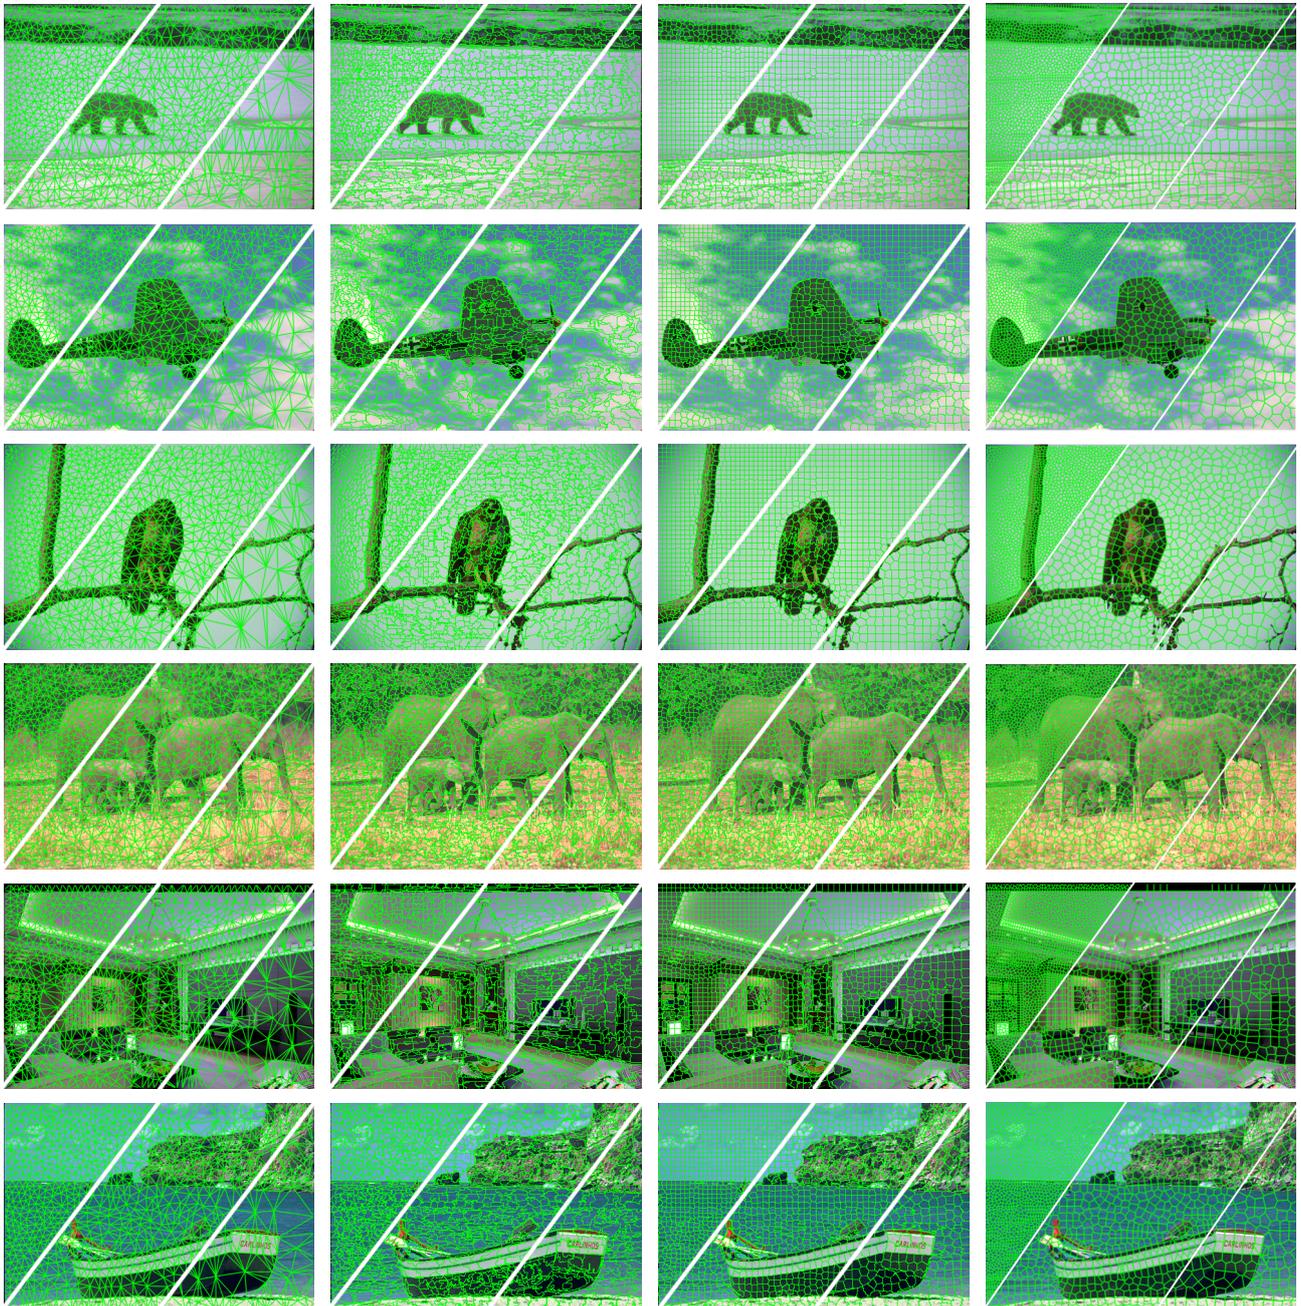


Figure 7. Visual comparison with oversegmentation methods. Although the results produced by SEEDS (second column), SLIC (third column) and Voronoi-based method (fourth column) provide accurate decompositions, region shape and size are rigid and poorly controllable. In contrast, our method (first column) offers simple, compact and flexible representations.

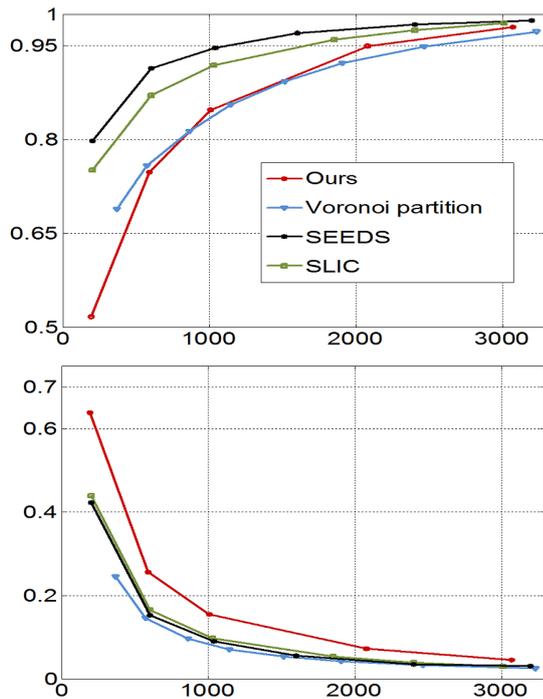


Figure 8. Quantitative evaluation on Berkeley database. Boundary recall (top) and undersegmentation error (bottom) are compared with three other methods: SEEDS, SLIC and Voronoi partition.

Descombes, X., 2011. Stochastic geometry for image analysis. John Wiley & Sons.

Duan, L. and Lafarge, F., 2015. Image partitioning into convex polygons. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3119–3127.

Green, P. J., 1995. Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika* 82(4), pp. 711–732.

Hastings, W. K., 1970. Monte carlo sampling methods using markov chains and their applications. pp. 3119–3127.

Labatut, P., Pons, J.-P. and Keriven, R., 2007. Efficient multi-view reconstruction of large-scale scenes using interest points, delaunay triangulation and graph cuts. In: 2007 IEEE 11th international conference on computer vision, IEEE, pp. 1–8.

Lee, D. and Schachter, B., 1980. Two algorithms for constructing a delaunay triangulation. *International Journal of Computer and Information Sciences* 9(3), pp. 219–242.

Levinshstein, A., Sminchisescu, C. and Dickinson, S., 2010. Optimal contour closure by superpixel grouping. In: K. Daniilidis, P. Maragos and N. Paragios (eds), *Computer Vision-Eccv 2010, Pt II, Lecture Notes in Computer Science*, Vol. 6312, pp. 480–493.

Liu, M., Tuzel, O., Ramalingam, S. and Chellappa, R., 2011. Entropy rate superpixel segmentation. In: 2011 IEEE Con-



Figure 9. Polygonal contouring. Using a basic edge tracking procedure as subsequent step to our method (see text), one can extract the contours of a specific class of objects (here building roofs) as low complexity polygons (green curves). Even if we note some errors, most of roofs are correctly polygonalized with such a basic tracking procedure.



Figure 10. Image abstraction. With only a few hundred triangles, one can generate an artistic representation of images that is both compact and detail-preserving (see the eyes of two birds on the left image).

ference on Computer Vision and Pattern Recognition, CVPR 2011, June 20, 2011 - June 25, 2011, pp. 2097–2104.

Martin, D., Fowlkes, C., Tal, D. and Malik, J., 2001. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001, Vol. 2, IEEE, pp. 416–423.

Mikolajczyk, K. and Schmid, C., 2005. A performance evaluation of local descriptors. *IEEE transactions on pattern analysis and machine intelligence* 27(10), pp. 1615–1630.

Moore, A. P., Prince, S. J., Warrell, J., Mohammed, U. and Jones, G., 2008. Superpixel lattices. In: 2008 IEEE conference on computer vision and pattern recognition, IEEE, pp. 1–8.

Mori, G., 2005. Guiding model search using segmentation. In: Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1, Vol. 2, IEEE, pp. 1417–1423.

Ren, X. and Malik, J., 2003. Learning a classification model for segmentation. In: Proceedings: Ninth IEEE International Conference on Computer Vision, October 13, 2003 - October 16, 2003, Vol. 1, pp. 10–17.

Ren, X., Fowlkes, C. C. and Malik, J., 2005. Scale-invariant contour completion using conditional random fields. In: Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1, Vol. 2, IEEE, pp. 1214–1221.

Sun, X., Christoudias, C. M. and Fua, P., 2014. Free-shape polygonal object localization. In: European Conference on Computer Vision, Springer, pp. 317–332.

Tuytelaars, T., 2010. Dense interest points. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, IEEE, pp. 2281–2288.

Van den Bergh, M., Boix, X., Roig, G., de Capitani, B. and Van Gool, L., 2012. Seeds: Superpixels extracted via energy-driven sampling. In: European conference on computer vision, Springer, pp. 13–26.

Veksler, O., Boykov, Y. and Mehrani, P., 2010. Superpixels and supervoxels in an energy optimization framework. In: K. Daniilidis, P. Maragos and N. Paragios (eds), Computer Vision-Eccv 2010, Pt V, Lecture Notes in Computer Science, Vol. 6315, pp. 211–224.

Zhang, Y., Hartley, R., Mashford, J. and Burn, S., 2011. Superpixels via pseudo-boolean optimization. In: 2011 IEEE International Conference on Computer Vision, ICCV 2011, November 6, 2011 - November 13, 2011, pp. 1387–1394.